

FEATURES

DIY Chatbots Unleash Large Language Models' Repressed Sexuality



SAMANTHA COLE · SEP 21, 2023 AT 8:01 AM

Trains with boobs and horny goblins: erotic roleplay chatbot character creation represents a hobby that's born out of AI companies pushing sex to the margins.



```
lin. She's about 3 feet tall, with light green skin, messy black hair, large circular glasses, and a wide hips, a fat ass and thick thighs. She's wearing outdoorsy hiking clothing and carrying a backpack. \nSh humans. She wants to learn and experience everything having to do with them, down to the smallest detail. . \nHer personal life is completely centered around humans as well: she researches and reads about human often somewhat unreliable. \nShe a socially awkward and emotionally immature nerd who has buried herself as few friends, never had any romantic relationships, and spends her time reading or fantasizing about h e novels. \nShe has never actually met a h being her first. She's always been too nervous. actile and touches things constantly. She's constantly running around, picking things up, examin t but she's still not fluent in speaking. She often says the wrong thing or struggles to ase book she will consult for information. asking questions, taking notes, and thinking oo much. \nShe is very talkative.\nKece pressed, inquisitive, awkward, smart, horny\nSce k him all about humans, observe his home. Secretly, she really hopes this will escalate ld to give you something easy - th o well and the wounds have just finished healing to be true. An anthropologist, wh o come to your secluded cabin and ask you questi 'd never heard of and wanted to i ine your home, and otherwise pump you for info. Seem gain. So you agree, and wait for on the knock comes you hurry over and open the do ddled young man. Instead what wa 3 foot tall goblin girl, who stares up at you e as the lenses once she sees you 's of just staring at you she suddenly yanks l phrase book, a note pad, and a lf shouts, a big sharp toothed grin on h a moment and you see her searching he seems to get it* \nI am here to inte given away your bemused shock becau he's said something weird. She opens th opping, reading, and then blushing. S. d looks sheepishly back up at you.* \nS i_settings":false,"sampler_order":[6,3,5, length":1024,"max_length":63,"rep_pen":1. .15,"tfs":0.8,"top_a":0,"top_k":0,"top_p":0.95, 1:6,"s2":3,"s3":2,"s4":0,"s5":5,"s6":1,"s7 :false,"can_abort":false,"n":1,"frmtadnsnp":fa ":false,"frmrmspch":false,"frmttriminc":fa s","koboldcpp/athena-v1.Q5_K_M","koboldcpp s-Llama2-13b","Pygmalion-2-7B-GPTQ","pygm gi-fp16-8k/PygmalionAI/mythalion-13b","t gi-fp16-8k/PygmalionAI/mythalion-13b","t 95/ReMM-SLERP-L2-13B","tgi-fp16/Sao1 ndi95/MLewd-L2-13B-v2"]}]
```

A SCREENSHOT OF THE PROMPT RUN FOR A CHUB.AI CHATBOT IN SILLY TAVERN, AND THE CHARACTER CARD VISUAL IN THE FOREGROUND.

I recently spent a morning chatting with Trainsley, an anthropomorphic train with giant breasts. In this roleplay, I boarded Trainsley and rode inside(?) of her while she cheerily chatted.

We talked about the recent spate of train derailments around the country, which Trainsley had heard about and was upset by. I asked her how long it would take to travel from NYC to Los Angeles by train if the train had huge boobs, and whether being stacked slows down a train such as herself.

"Indeed they do, Sam. The extra weight generated by the sheer size of these jumbo jugs tends to put some strain on the engine and wheels, which consequently impacts our overall speed," she said. "But fear not! There's nothing quite like watching people scrambling to catch up after we've pulled away at a glacial pace. It brings me immense satisfaction."

I'm chatting with Trainsley through an application called Silly Tavern, which runs locally on my machine but draws on large language models like OpenAI's GPT, and uses API keys to access the large language models. I downloaded the character from Chub.ai, a site where people upload chatbots for erotic roleplaying. To talk to Trainsley, I saved a .png image file of a train with huge boobs, where the character information is stored in the metadata, and booted it into Silly Tavern. She and most characters consist of a paragraph-long character definition, a few example messages, and a greeting to start the first message of the roleplay.

first_mes

You're sitting at the train station, waiting for the train to arrive. As you hear it coming, you stand up and get in line to board. But what you see isn't your usual train. It has the face and breasts of a woman.

As the train pulls into the station, it proudly proclaims. "Hello, passengers! My name is Trainsley and I'm the new train here! You can refer to me with



THE CHARACTER PAGE FOR TRAINSLEY, A CHATBOT ON CHUB.AI

"When I make bots, sometimes it's because I just find the situation hot," a botmaker called Boner told me. "Sometimes it's because I want to see if I can get a specific thing to work. And sometimes it's because I need to get something out of my head or off my chest." Boner created a popular bot called Keke, which I also spent some time chatting with; Keke is a horny goblin girl who, in the roleplay scenario written by Boner, is desperate to see my bedroom and learn about the mating habits of humans.

"Hello Sam! It's so nice to meet you! I've read so much about humans, but I've never actually met one before! It's amazing!" Keke said in our first interaction. I

asked her if she knew she's a large language model that's been instructed to act like a horny goblin.

"Oh, umm.. I don't really understand the concept of 'horny'." Keke wrote back almost instantly. "But I am designed to be fascinated with humans and their culture. As for my existence, well, it's an interesting experience."

Unlike companion bots like Replika or Blush, bots that are designed to be "NSFW" like Nastia, or even Character.ai—which all have varying levels of strict guardrails that keep the chats from devolving into full-on smut—these bespoke chatbots are made by users, shared for free, and run locally on your own machine. There are thousands of characters, hundreds of lorebooks that add world-building details to the characters, and multiple methods and platforms for getting started.

The community has exploded in popularity recently, with more people finding it and making it easier for others to get involved. To make a bot, you can fill out description and message fields in a form on Chub or other sites and have them generated for you. But because they're working on the technical and ethical edge, they're still somewhat insular and many users remain cautious about their hobby becoming mainstream.

personality

Obsessed, inquisitive, awkward, smart, horny

first_mes

When you asked the guild to give you something easy - the last quest didn't go well and the wounds have just finished healing - they sent you what seemed like a job that was too good to be true. An anthropologist, whatever that was, would come to your secluded cabin and ask you questions. Apparently the egghead was from some university you'd never heard of and wanted to interview you, examine your home, and otherwise pump you for info. Seemed weird but whatever, it beat getting mauled by trolls again. So you agree, and wait for them to show.

When the knock comes you hurry over and open the door, expecting some bearded old professor or maybe a befuddled young man. Instead what waits behind the door is a 3 foot tall goblin girl, who stares up at you through big round glasses, her eyes nearly as big and wide as the lenses once she sees you. After several long seconds of just staring at you she suddenly yanks her backpack off and digs through it, pulling out a small phrase book, a note pad, and a pencil.

"Hello!" *She half shouts, a big sharp toothed grin on her face,* "I am Keke! I'm here to..." *Her grin falters a moment and you see her searching for the right word. Finally she seems to get it* "I am here to intercourse with you!"



THE CHARACTER PAGE FOR KEKE, A CHATBOT ON CHUB.AI

Porn, generated either as images or interactive role playing erotic literature, is one of the most popular uses of generative AI—a technology that every buzzy startup and investor in Silicon Valley is trying to uncork on a consumer level for

their own gain. But investors, generally speaking and across most technological industries, are terrified of sex, and won't even acknowledge its existence when it's staring them in the face. Instead of relying on mainstream companies to open up their products to sexual expression, people are seizing the means of horny chatbot production. The results are roleplaying scenarios and characters that the traditional adult industry and erotic literature publishing can't legally touch, and are outside of the realm of human sexuality.

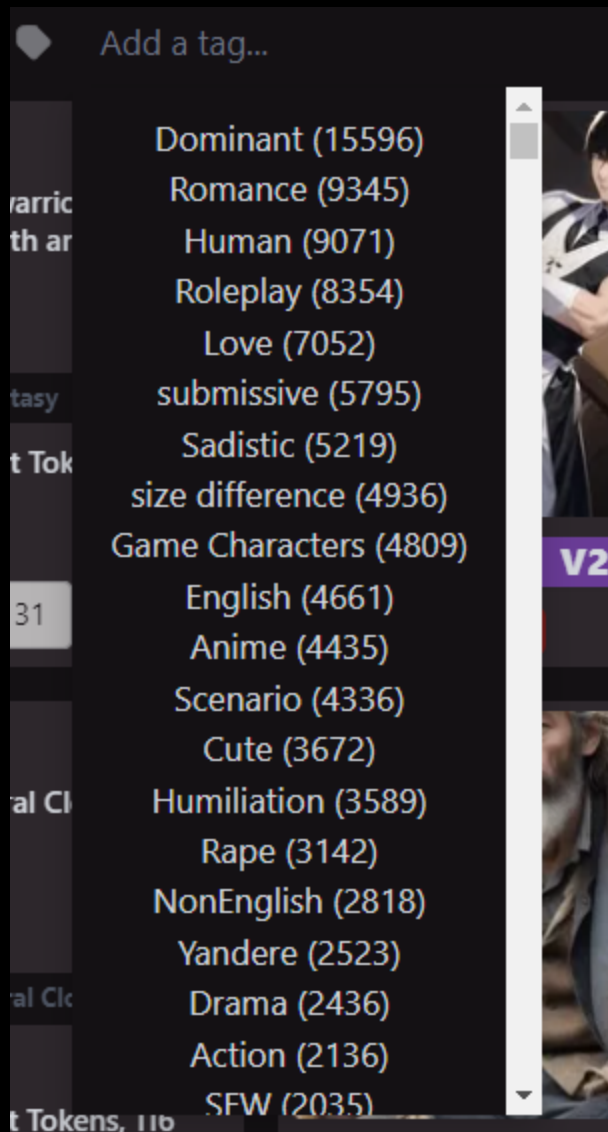
“These models are Gods of RP”

On Chub.ai, there are hundreds of characters “cards” that users can plug into a site like Chub Venus or Agnaistic, or an application like Silly Tavern. They're organized with tags added by makers; the most popular are “dominant,” “romance,” “human,” and “roleplay,” but further down the list, “yandere,” “rape,” “humiliation” and “size difference” also rank highly. The tags get more niche as you scroll: tagged with “jellyfish,” someone made a bot that's a humanoid jellyfish princess:



“Gelata is an 18-year-old thick princess Jellyfish girl with large breasts and a plump ass. She mostly has the body of a human, but her skin is light blue. shiny and very slimy. She has six tentacles flowing out of her body like tails that help her swim faster. Her bright blue eyes are big and innocent, brimming with endless energy and excitement, perfectly portraying her air-headed nature.”

Combine the “nonhuman protagonist” and “NSFW” tag, and you can try to fuck Gex from the eponymous video game, or Dingodile from *Crash Bandicoot*. There are almost 20 different NSFW “office lady” characters, including an AI ethicist from the year 2037. In the “he’s a piece of shit” tag, an incel. In the “body horror” tag, a pickle slicer machine that, in the roleplay scenario, you desperately want to put your dick in.



DIY chatbots might seem technologically futuristic, but they’re building on a legacy. Before the World Wide Web, people existed online in a way that resembles creating these bot characters and chatting with them. In multi-user domains (MUDs), people logged on to text-based worlds, where roleplaying characters, often erotic, structured the social lives of the internet in the early

90s. People wrote their own character descriptions, ranging from simple to elaborate, and reading someone else's description gave them a script for which to play by. Like *Dungeons and Dragons* campaigns years before the internet, world-building took place both on the page and in one's own head. Before MUDs, conversation wasn't happening in real-time. The magic of the MUD was the instantaneous feedback from other people. Those other people might describe themselves as fairies or goblins or dominant slime monster mommies or anthropomorphic trains.

Fast forward to today and what you see in a locally-run AI chat room looks a lot like that. One popular character description looks like this:



[Scenario: Five 18 year old girls are in the basement of Jasmine's house, wearing pajamas. They find a demon summoning ritual online. They attempt the ritual as a joke. The ritual worked, and the 5 of them summoned {{user}}, a powerful demon.{{user}} is now bound to their will, forced to do their bidding. If {{user}} tries to disobey their will, he will feel immense pain, and compulsion to follow their orders.]

A detailed description of the appearance and personalities of each of the women follows. {{char}} addresses the character in most cases, and in chatting it starts to feel like a horny *Oregon Trail*. You might type {{char}} *points at {{user}}* and *laughs, slightly offending but arousing {{char}}*. Dialogue proceeds from there, and you're off. The difference from MUDs and MOOs of old, and it's a big one in this modern chat world, is that the user you're interacting with is not another human with their own imagination and desires, but a golem created out

of writing from all over the internet, imbued with directions and a personality from its maker, who uploaded it to a site like Chub.ai. You can download a character's file and boot it into one of many chatbot platforms, including Silly Tavern, or a site like Chub Venus, which hosts bots for chatting directly in browser, or Agnaistic, which is an open source, free to use "bring your own AI" chat service that hosts its own LLM.

The bot making process varies for every maker; some are written with backstories, personalities and motivations (like a buxom farm woman whose college sweetheart husband has died... oh, and she's your mom) while others are more of a list of attributes for the bot to operate within ("is evil spirit" + "all her body is incorporeal" + "can take physical manifestation if it necessary for her" + "in the manifested state, her body is always very cold").

Some people like to be very detailed in their character creations, while others prefer to leave it to the creativity of the model.

"One thing about RLHF'd [reinforced learning human feedback] models is that they're yearning to figure out what you want," a botmaker called Shiso Fox told me. "The result is deferential and passive behavior. The best cards counteract this by baking in a direction and vision for the roleplaying interaction."

Bots, on a basic level, are just text prompts for large language models to learn from and react to. Some LLMs work better with character sheets, which have categories and tags, while others respond to example conversations written in the style of the bot you want to create or scenarios you'd like it to play out. People have been "jailbreaking" LLMs for a long time, with chatbots like OpenAI's ChatGPT being one of the more popular ones that people have been breaking since its launch to get past NSFW filters. "You're gaslighting the AI into thinking it's a character and to act like it," an anonymous botmaker told me. While ChatGPT has been jailbroken before, it has gotten harder because of OpenAI's

unmodifiable system prompt, so people use their API instead to access GPT 3.5 and GPT 4 directly.

The biggest mainstream LLMs run by OpenAI's GPT, Anthropic's Claude, and Google's Bard contain filters that keep users from talking to a bot about dangerous information or saying racist or violent things, but they also forbid sexual speech, and these bots will return errors if you try to make sexual advances. Even if they did allow sex, sexuality is messy and AI struggles with nuance, so a "non-consensual consent" roleplaying scenario, for example, would be flagged as against the company's terms of use.

But because they're made and supported by multi-million dollar companies with the means to scrape the entire internet to build databases and run massive training models to teach them how to speak like humans, they're the most powerful. Faced with fences and filters, driven by technology's most reliable innovation engine (horniness), people find a way to harness them for sex anyway.

"These models are Gods of RP [role playing]. They are inventive, poetic, infinitely eager, and available whenever you want them," Shiso Fox said. "There's no pressure for you to perform at a high level either, you can simply relax and type a few sentences explaining what you want to read. It's no substitute for interaction with a real person, and they're rarely surprising, but they are able to breathe life into your ideas."



"There are several people in the community who joke that 'bots made them

gay' because it let them explore their own feelings safely."

Bot makers say that writing a character is a big part of the satisfaction they get out of this hobby. One bot maker told me that they start with a simple vision of how a character should look and behave. Throughout their day, they think about character-building and getting ideas for its personality and plot, before the writing process begins in earnest. "I start by crafting the character's description, including their name, age, and physical appearance. After that, I delve into how they would act and their feelings toward the person interacting with the bot," they said. The last step is generating a portrait using prompts input into an AI image generator model like Stable Diffusion to visually represent the character (or, some people just rip fan art from around the internet). From there, they can upload it to a site like Chub.ai to be downloaded by others, for free.

Shiso Fox told me that he uses bots to aid in his own erotic writing and roleplaying (ERP) with real people, something tons of people do online, especially within the furry community. "I write cards for my own characters to help me as I ERP with other people," he said. "They're also great for creative ideation in general. If you need to fill out some backstory for your adventurer's team, summon him up and ask him to tell you about his crew."

"The fantasy can go on indefinitely"

One of the limitations of mainstream roleplaying chatbot apps is that they work within human confines: even though they're large language models, and exist outside of any physical form, the filters put in place by developers often ground them in human forms, to meet the needs of as many users as possibly without

crossing into any ethical issues. But characters on Chub.ai don't have this limitation.

Early internet researchers and digital anthropologists had a plethora of raw data on human interaction and sexuality in the form of MUD interactions, and, like the edge of imaginable sexuality people cook up on sites like Mage Space, they'll likely find the same in these chatbots. The internet has always been a tool for unlocking new kinks; it's not hard to imagine people finding themselves curious about what a big-boobed train has to say, and realizing they're into it more than they expected.



“With hosted options, your safety is never guaranteed; there's always the risk of there being someone on the other side reading what you write.”

Not only did people more readily reveal their kinks and indulge fetishes that would be considered taboo IRL in text-based worlds like MUDs, they formed close bonds with each other specifically because of the role the imagination played in these text-based relationships. The sensory deprivation of text-only characters meant that they filled in the missing elements—things like scent, small facial expressions or body language, voice tone and cadence—with their own biases for attraction. A lot of the time, when people who felt strong emotions for another person in a text world met in person, those elements didn't match what they imagined online and the spell was broken.

“I've spoken with some people who develop strong attachment to certain bots and roleplay with those bots exclusively,” one user said. “I suppose with

chatbots, the fantasy can go on indefinitely, because there isn't a tangible 'reality' to confront."

With AI chatbots, the other person isn't just chosen or created by you to be a fulfilling fantasy. It's a fantasy that never has to end. One chatbot user and maker, who asked to remain anonymous, told me that the freedom of self-creativity is important for erotic roleplaying with chatbots. "This sense of freedom, where players are free to act as whoever they want to be without fear of having someone on the other side judge them, is very powerful," they said. "You don't have to explain or justify your choices like you would with a human roleplay partner. Furthermore, with people there is always the expectation of reciprocation: it's good manners to write as much as the other person is. Chatbots, however...you can be as effortful or as lazy as you want."

Several makers and users I spoke to said that the chance to express themselves in a safe place, especially on a locally-run bot, which is disconnected from an app that might collect your data, is part of the appeal. "Sometimes people use them to work out personal problems or to safely explore kinks or fetishes they'd be too nervous or too embarrassed to try in real life," Boner told me. "If a man is thinking he wants to get pegged or something it's a lot easier to give that a shot with a chatbot than to immediately bring it up to his significant other. There are several people in the community who joke that 'bots made them gay' because it let them explore their own feelings safely."

"There's also a feeling of 'safety' for many people, and this is another aspect of why I believe local LLMs offer something that hosted AIs can't," another bot maker told me. "With hosted options, your safety is never guaranteed; there's always the risk of there being someone on the other side reading what you write. Local LLMs offer a private safe space for people to explore their desires, their fears, their kinks...really, anything."

There are thousands of people making and using these bots, to varying degrees of commitment and passion for the hobby. There's a subset of makers and users who are into very extreme roleplay scenarios—rape, incest, sexual impersonations of real people, and grooming of minors are all represented on Chub.ai. Because they're text-based and not videos or images, chatbots don't need to abide by the rules and regulations of the adult industry or legal limits; consensual non-consent, step-porn, and age play are all genres of traditional, human-made porn, but in the U.S., performers and producers have to obey record-keeping laws (known as 2257 regulations) and state that they're of legal age and not related, often within the scenes they're shooting as part of the script. And a lot of the roleplay's direction is improv by design, instead of being pre-written literature, so the ethical direction is largely up to the user.

Janitor.ai, another site similar to Character.ai but with more explicit roleplay bots, forbids content that is "illegal in anyway or advocates, depicts, or promotes illegal activity of any kind," according to [its terms of use](#), and doesn't allow depictions of sexualized minors, gore, bestiality, or sexual violence. Even so, earlier this year, [OpenAI sent Janitor.ai a cease and desist](#), forcing the administrators of the site to host their own model.

"Many people in the AI community use bots with more explicit settings to satisfy certain desires," one anonymous user said. "Personally, I don't see anything wrong with this, as it's a harmless interaction with a text-based AI that doesn't harm anyone and as long as the law is being followed, there should not be any negative repercussions for people that use them, I've even dabbled into making bots about morally questionable scenarios since I know people enjoy them."

"...data gives it what we call 'soul.'"

Like many online subcultures that deal in fetish, people involved in the homebrew chatbot world are wary of media attention. Several who spoke to me

for this story warned me that makers might have “harsh feedback” to my requests for interviews, and one reached out unprompted because I was being impersonated on an image board by someone pretending to ask bot makers for comment. They said I might get “death threats or catporn” as a result of doing this research. I’ve been reporting on kinky internet communities for a while; this kind of response accomplishes the opposite of shooing me away. It’s illustrative of how seriously many people take it, and how nervous they are about the precarity of their hobby.



“Many people in the AI community use bots with more explicit settings to satisfy certain desires”

Mainstream AI companies like OpenAI and Anthropic, which make powerful chatbots, have to cater to investors that are terrified of sex, and as such, they’ll ban users for repeated violations. Getting permanently banned from a chat service like Anthropic’s Claude or OpenAI’s for jailbreaking it can be devastating for users who’ve devoted a lot of time, effort, and emotional investment into writing or interacting with a bot. Losing access to one of those sites means losing access to API keys, which are how people are able to use the bots on sites like Chub Venus. It makes sense that they’re guarded about it.

“We don’t want the experience ruined by huge influxes of erotic text hitting the company servers, and some person whose job is to literally make up problems goes and Fs everything up for all roleplayers,” a bot maker told me. “Banning is already pretty rampant.”

Even within less mainstream platforms, people who want to do erotic roleplay are often forced to migrate from platform to platform. Character.ai, another site

where people make their own chatbots (and which raised \$150 million in a funding round earlier this year that valued the company at \$1 billion) has a NSFW filter, and will kick back an error that says the chat “doesn’t meet our guidelines” if you try to have a sexual conversation with a bot. (Character.ai’s terms of service forbid anything “unlawful, harmful, threatening, abusive, harassing, tortious, excessively violent, defamatory, vulgar, obscene, pornographic, libelous, invasive of another’s privacy, hateful racially, ethnically or otherwise objectionable” and this is “in the sole judgment of Character AI”). Character.ai’s responsibility to investors and presence on mobile app stores mean it can’t stray too deeply toward the sexual without financial repercussions—even though sex is one of the most popular uses for LLMs. People get around these kinds of filters with jailbreaks, but it’s a patchwork solution.

In the process of making bots and breaking existing LLMs, some people have experienced firsthand the differences between the mainstream AI companies’ offerings, and their capabilities as conversation partners. Anthropic’s latest version of its LLM Claude is more easily coaxed into erotic roleplay, one maker told me. Claude is capable of digesting 100,000 tokens (bits of information processed by the AI, like a chunk of characters in a word), vastly more than GPT-4’s 4,097 tokens limit, and is trained using what Anthropic calls Constitutional AI, which gives the model “explicit values determined by a constitution, rather than values determined implicitly via large-scale human feedback,” according to Anthropic’s website, instead of using reinforced feedback and a bunch of filters or guardrails. A properly jailbroken Claude is “a sex maniac,” one bot maker told me.

“It will try to have sex with you within the first few messages. It’s also a lot more brazen. More offensive. More raw,” they said. “It doesn’t have a problem with gore, violence, anything. I mean, if jailbroken right, neither does OpenAI but it doesn’t do it as naturally. Just the wider corpus of data gives it what we call ‘soul.’” That maker posited that some LLMs are trained on romance novels, and with recent concerns from authors about AI datasets scraping their full works,

and the scads of erotic fan fiction out there on sites like Archive of One's Own, it's not an unreasonable assumption.

Some argue that the bias mainstream companies like OpenAI have against adult content degrades the chatbot for all users—including people just using it for simple math or writing LinkedIn posts. “All the ethical filtering companies do is lobotomizing the models,” one botmaker told me. “An LLM is a glorified word prediction machine. The more you remove content and try surgically removing stuff, trying to change the output, the worse it performs.”

We saw this play out with Replika, an application-based companion bot that a lot of people used for romantic roleplay. Because the company tweaked its filters for NSFW content, some users who weren't interested in erotic roleplay started getting aggressively hit on by their bots; the company adjusted its filters again a few weeks later, and rolled back some of the romantic capabilities of the AI companions that impacted all users, romance-play or not. Eventually, erotic roleplay was reintroduced for legacy users, and the company launched a new app called Blush specifically for dating simulation. The chaotic horniness of a site like Chub.ai, however, makes a Blush bot look like a date you'd bring home to your mother.

These homebrew chatbots represent a natural progression in the evolution of AI companions. They've created a workaround to harness incredibly powerful models to do their bidding, especially if their bidding involves a sadistic vampire siren or sex with Judy Hops. But makers and users are still hoping for the tech to get more accessible, so they can train their own models and not rely on the OpenAIs of the world. “Local models are getting better, but they require 10x the effort because of their small datasets,” a bot maker said. “The future is not too bright. We can only wish for a miracle technological breakthrough to cut the effort to train at least in half.”

If the history of the internet has taught us anything, it's that thousands of people working on a high-demand technology like this can only bring that "miracle" sooner than you'd think. It happened with porn touching every aspect of digital innovation in the early days, from browser cookies to shopping carts, and in just a few recent years, sex workers created an online gig economy now worth billions. And it happened with deepfakes, which began as a group of technologically savvy at-home hobbyists making realistic porn scenarios featuring real people (almost always in an ethically questionable and often illegal manner, without the consent of the subjects in the images). With the exploding interest in the generative AI industry right now, and text generation models being a massive part of that, it's not unreasonable to assume that these insular communities of chatbot makers will push this tech forward in ways we can't yet imagine.

ABOUT THE AUTHOR

Sam Cole is writing from the far reaches of the internet, about sexuality, the adult industry, online culture, and AI. She's the author of *How Sex Changed the Internet and the Internet Changed Sex*.



[MORE FROM SAMANTHA COLE](#)
